

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s): Ashok Singhal, David J. Broniarczyk, George R. Cameron, Jeff A. Price

Assignee: 3PAR Inc.

Amended Title: Node Controller for a Data Storage System

Serial No.: 09/883,681

Filing Date: June 19, 2001

Examiner: Nimesh G. Patel

Group Art Unit: 2112

Docket No.: M-8496 US
(3PD-M-8496 US)

Confirmation No. 1044

San Jose, California
September 17, 2008

Mail Stop Appeal Brief - Patents
Under Secretary of Commerce for Intellectual Property and Director of the
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, Virginia 22313-1450

APPEAL BRIEF UNDER 37 CFR §41.37(c)

Dear Sir:

Applicant hereby appeals the rejections of claims 9, 13, 16, 18, 19, 21 to 25, 29, and 35 to 61 in the March 17, 2008 Final Office Action to the Board of Patent Appeals and Interferences. A one-month extension is hereby requested, extending the period for response to September 17, 2008.

Please charge the fees of \$510.00 for this Appeal Brief as set forth in 37 C.F.R. §41.20(b)(2), and \$120.00 for the extension of time set forth in 37 C.F.R. § 1.17(a)(1), to the authorized credit card. Please also charge any amounts underpaid or credit any amounts overpaid to Deposit Account No. 502226.

REAL PARTY IN INTEREST

The real party in interest is the assignee 3PAR, Inc.

RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

STATUS OF CLAIMS

Claims 9, 13, 16, 18, 19, 21 to 25, 29, and 35 to 61 are pending, twice rejected, and appealed. Claims 1 to 8, 10 to 12, 14, 15, 17, 20, 26 to 28, and 30 to 34 have been canceled.

STATUS OF AMENDMENTS

The Examiner issued the Final Office Action on March 17, 2008. Applicant did not file any amendments subsequent to the rejection.

SUMMARY OF CLAIMED SUBJECT MATTER

Claim 18

Claim 18 recites a node controller for a node in a data storage system having at least two nodes. Specification, p. 5, line 25 to p. 6, line 14; p. 7, lines 1 to 6; Fig. 1 (e.g., a node 10 and a node controller 12). Each node includes one computer-memory complex and one node controller distinct from said one computer-memory complex. *Id.*, p. 7, line 25 to p. 8, line 3; p. 8, lines 20 to 22; Fig. 1 (e.g., computer-memory complex 18 and node controller 12). The node controller includes a memory controller, an input/output bus interface, and a logic engine. *Id.*, p. 13, lines 9 to 19; p. 14, lines 11 to 17; p. 15, lines 13 to 24; Fig. 3 (e.g., node controller 12, memory controller 54, I/O bus interface 50a and 50b, and logic engines 56).

The memory controller is operable to interface the node controller with (1) a cluster memory that stores data being transferred through the node, and (2) a link coupled to another node controller in another node of the data storage system. *Id.*, p. 14, lines 11 to 23; p. 9, lines 6 to 11; Fig. 3 (e.g., memory controller 54, cluster memory 20, and communication paths 16).

The input/output bus interface is operable to interface the node controller with an input/output bus coupled to the computer-memory complex of the node, at least one host device and at least one storage device. Id., p. 13, lines 11 to 19; p. 11, lines 10 to 24; Figs. 2 and 3 (e.g., I/O bus interface 50a and 50b, node controller 12, I/O buses 36a and 36b, and computer-memory complex 18)

The logic engine is coupled to (1) the memory controller, (2) the link, and (3) the input/output bus interface. Id., p. 15, lines 13 to 15; Fig. 3 (e.g., logic engines 56, memory controller 54, communication paths 16, and I/O bus interface 50a and 50b). In a first type of data transfer, the logic engine performs a logic operation to a plurality of data from one of a plurality of data sources in the data storage system and writes the result of the logic operation to one of a plurality of data destinations in the data storage system. Id., p. 15, lines 19 to 24. The data sources include the cluster memory and the input/output bus interface, the data destinations comprising the cluster memory, the link, and the input/output bus interface. Id., p. 15, lines 19 to 24; Fig. 3 (e.g., cluster memory 20, I/O bus interface 50a and 50b, and communication paths 16). The logic operation is used to calculate a parity data for writing a full or a partial RAID stripe or to reconstruct a lost data using the parity data. Id., p. 16, line 12 to p. 17, line 3.

Claims 9, 13, 16, 19, and 21 to 25

Claims 9, 13, 16, 19, and 21 to 25 depend from claim 18 and they are patentable for at least the same reasons as claim 18.

Claim 29

Claim 29 recites a node in a data storage system including at least two nodes. Specification, p. 5, line 25 to p. 6, line 14; p. 7, lines 1 to 6; Fig. 1 (e.g., a node 10 and a node controller 12). The node includes an input/output bus, a computer-memory complex, and a node controller. Id., p. 7, line 25 to p. 8, line 3; p. 8, lines 20 to 22; p. 11, lines 10 to 24; Figs. 1 and 2 (e.g., I/O buses 36a and 36b, computer-memory complex 18, node controller 12).

The input/output bus is coupled to at least one host and at least one storage device. Id., p. 11, lines 10 to 24; Fig. 2 (e.g., I/O buses 36a and 36b).

The computer-memory complex includes a central processing unit (CPU), a system memory storing information for controlling data transfer through the node, and a controller coupling the

CPU, the system memory, and the input/output bus. Id., p. 10, line 8 to p. 11, line 5; Fig. 2 (e.g., computer-memory complex 18, CPUs 30a and 30b, system memory 32, main controller 24).

The node controller is distinct from the computer-memory complex. Id., p. 8, lines 20 to 22; Fig. 1 (e.g., computer-memory complex 18 and node controller 12). The node controller includes a memory controller, an input/output bus interface, and a logic engine. Id., p. 13, lines 9 to 19; p. 14, lines 11 to 17; p. 15, lines 13 to 24; Fig. 3 (e.g., node controller 12, memory controller 54, I/O bus interface 50a and 50b, and logic engines 56). The memory controller is operable to interface the node controller with a cluster memory of the node that stores data being transferred through the node. Id., p. 14, lines 11 to 23; p. 9, lines 6 to 11; Fig. 3 (e.g., memory controller 54 and cluster memory 20). The input/output bus interface is operable to interface the node controller with the input/output bus. Id., p. 13, lines 11 to 19; Figs. 2 and 3 (e.g., I/O bus interface 50a and 50b, node controller 12, and I/O buses 36a and 36b). The logic engine is coupled to the memory controller, the input/output bus interface, and a link to another node of the data storage system. Id., p. 15, lines 13 to 15; Fig. 3 (e.g., logic engines 56, memory controller 54, I/O bus interface 50a and 50b, and communication paths 16).

At least one of the host device and the storage device coupled to the input/output bus is able to read and write the cluster memory via the input/output bus. Id., p. 14, lines 6 to 10; Figs. 2 and 3 (e.g., I/O buses 36a and 36b and cluster memory 20). The logic engine is able to transfer data from the cluster memory to said another node via the link. Id., p. 17, lines 8 to 10; Fig. 3 (e.g., logic engines 56, cluster memory 20, and communication paths 16). The memory controller is able to receive data from said another node via the link. Id., p. 17, lines 8 to 10; Fig. 3 (e.g., memory controller 54 and communication paths 16).

Claims 35 to 61

Claims 35 to 47 depend from claim 29 and they are patentable for at least the same reasons as claim 29.

Claim 48

Claim 48 recites a node in a data storage system having at least one node for providing access to a data storage facility. Specification, p. 5, line 25 to p. 6, line 14; p. 7, lines 1 to 6; Fig. 1 (e.g., a node 10 and a node controller 12). The node includes a node controller and a computer-

memory complex. Id., p. 7, line 25 to p. 8, line 3; Fig. 1 (e.g., computer-memory complex 18 and node controller 12).

The node controller is distinct from the computer-memory complex. Id., p. 8, lines 20 to 22; Fig. 1 (e.g., computer-memory complex 18 and node controller 12). The node controller includes an input/output bus interface, a memory controller, and one or more logic engines. Id., p. 13, lines 9 to 19; p. 14, lines 11 to 17; p. 15, lines 13 to 24; Fig. 3 (e.g., node controller 12, I/O bus interface 50a and 50b, memory controller 54, and logic engines 56). The input/output bus interface is operable to interface the node controller with an input/output bus that is coupled to the computer-memory complex, at least one host device, and at least one storage device. Id., p. 13, lines 11 to 19; p. 11, lines 10 to 24; Figs. 2 and 3 (e.g., I/O bus interface 50a and 50b, node controller 12, I/O buses 36a and 36b, and computer-memory complex 18). The memory controller is operable to interface the node controller with a cache memory in the node. Id., p. 14, lines 11 to 23; p. 9, lines 6 to 11; Fig. 3 (e.g., memory controller 54 and cluster memory 20). The one or more logic engines are coupled to the input/output bus interface, to the memory controller, and to a link that can be coupled to another node in the data storage system. Id., p. 15, lines 13 to 15; Fig. 3 (e.g., logic engines 56, I/O bus interface 50a and 50b, memory controller 54, and communication paths 16).

The node controller is arranged so that the computer-memory complex, the host device, and the storage device are able to access the node controller via the input/output bus so that the host device and the storage device are able to read and write the cache memory via the input/output bus. Id., p. 14, lines 6 to 10; Figs. 2 and 3 (e.g., node controller 12, computer-memory complex 18, I/O buses 36a and 36b and cluster memory 20). The logic engine is able to transfer data from the cache memory to another node in the data storage system via the link, and the memory controller is able to receive data from the another node in the data storage system via the link. Id., p. 17, lines 8 to 10; Fig. 3 (e.g., logic engines 56, cluster memory 20, communication paths 16, and memory controller 54).

Claims 49 to 61

Claims 49 to 61 depend from claim 48 and they are patentable for at least the same reasons as claim 48.

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The Examiner rejected claims 13, 18, 19, 21 to 25, 29, 35, 38 to 50, 52 to 54, and 56 to 61 under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,768,623 (“Judd et al.”). The Examiner rejected claims 9, 16, 36, 37, 51, and 55 under 35 U.S.C. §103(a) as being unpatentable over Judd et al. in view of what is well known in the art.

ARGUMENTS

Overview of the Claimed Invention

The claimed invention is directed to a node controller of a node in a data storage system. The node controller is distinct from a computer-memory complex of the node. The node controller is coupled to a cluster memory and to a high speed link to another node in the data storage system. The node controller is further coupled by an input/output bus to the computer-memory complex and at least one of a host device and a storage device.

The node controller is the component of the node responsible for data transfer to and from the link. In contrast, the computer-memory complex is responsible for functions related to the control of the node. This split construction has the advantage that the computer-memory complex does not need to be burdened with data transfers and can perform its control functions independent of the data transfer load of the node. This is because the host or storage device may read and write the cluster memory so that the system memory of the computer-memory complex is not burdened with temporarily storing the data being transferred through the node. In addition, the node controller moves the data to another node through the link so that the input/output bus is not burdened with data transfer between nodes. Furthermore, the node controller includes a logic engine can be programmed to write RAID stripes and reconstructs lost data from RAID stripes so that the processor of the computer-memory complex is not burdened with the RAID operations.

Claim 18

The Examiner cited various elements of Judd et al. against the elements of claim 18. Applicant summarizes the correspondence in the table below.

Claim term	Judd et al., Fig. 7 unless otherwise indicated
Node	Hosts 80-81, adapters 86-87
At least two nodes	Node 1: Hosts 80-81, adapters 86-87 Node 2: Hosts 82-83, disk arrays 94-95
Computer-memory complex	Host 80
Node controller	Host 81
Memory controller	Cache controller 68 (Fig. 6)
Cluster memory	Cache controller 68 (Fig. 6)
Link to another node controller	SSA loop 98 to host 82
Input/output bus interface	Adapter 87
At least one host device	Adapter 86 (note that adapter 86 is actually part of host 80)
At least one storage device	Disk array 92
Logic engine	XOR hardware 42 (Fig. 5); RAID controller 70 (Fig. 6)

In the March 30, 2007 Amendment, Applicant submitted that the Examiner arbitrarily combined hosts 80 and 81 of Judd et al. as the recited node of claim 18. In the March 17, 2008 Final Office Action, the Examiner responded that a “node can comprise any number of devices” and the combination of hosts 80 and 81 is reasonable. March 17, 2008 Final Office Action, p. 5. The Examiner did not provide any further reasoning or support. Applicant submits again that it is unreasonable to cite the combination of hosts 80 and 81 of Judd et al. as the recited node.

Judd et al. never refers to the combination of hosts 80 and 81 as a single node. Instead, Judd et al. clearly shows hosts 80 and 81 as individual nodes in SSA loop 98 that connect all of hosts 80

to 85. Hosts 80 and 81, like hosts 82 to 85, are individual nodes because they are identical to each other in hardware and software, and they serve the same function of providing primary and backup access to their respective disk arrays.

Claim 18 specifically recite a node controller that is distinct from a computer-memory complex located in the same node as the node controller. As discussed above, this split construction has the advantage that the computer-memory complex does not need to be burdened with data transfers and can perform its control functions independent of the data transfer load of the node. The Examiner, however, has cited two identical hosts 80 and 81 against the recited node controller and the recited computer-memory complex that are very different. Applicant respectfully submits that it is unreasonable to read the recited node with two different elements with different functions onto two identical elements that serve the same function.

In the March 30, 2007 Amendment, Applicant also argued that the Examiner cannot cite adapter 86 in host 80 as the recited host of claim 18 because the Examiner had already cited host 80 against the recited node of claim 18 (a separate and distinct claim limitation) and adapter 86 is part of host 80. In the March 17, 2008 Final Office Action, the Examiner responded that the “specification describes a host device as a processing device that may include various peripheral devices (Page 6, Lines 3-7)” and “Judd clearly shows adapter 86 as a peripheral device.” March 17, 2008 Final Office Action, p. 5. Applicant submits again that the Examiner has impermissibly cited one element in Judd et al. against two separate and distinct elements in claim 18.

While a host device may certainly include peripheral devices, it does not make a host device a peripheral device of a node. Claim 18 clearly recites the node and the host as distinct and separate elements. As the Examiner already cited hosts 80 and 81 as the recited node, he cannot cite a component of host 80, specifically adapter 86, as the recited host of claim 18.

Applicant respectfully submits that when interpreted in a reasonable manner that is consistent with its specification and common sense, Judd et al. does not disclose a node controller distinct from a computer-memory complex located in the same node as the node controller. This is the reason why that, even assuming the interpretations proposed by the Examiner in the last three office actions are correct, these interpretations are always missing elements recited in claim 18.

For the above reasons, amended claim 18 is patentable over the cited references.

Claims 9, 13, 16, 19, and 21 to 25

Claims 9, 13, 16, 19, and 21 to 25 depend from amended claim 18. Thus, claims 9, 13, 16, 19, and 21 to 25 are patentable for at least the same reasons as amended claim 18.

Claim 29

Claim 29 recites a node having “an input/output bus coupled to at least one host and at least one storage device.” As discussed above with regards to claim 18, Judd et al. does not disclose the recited host if the Examiner cites hosts 80 and 81 as the recited node. Thus, amended claim 29 is patentable over Judd et al. for at least the same reasons as amended claim 18.

Claims 35 to 61

Claims 35 to 47 depend from amended claim 29 and are patentable over the cited references for at least the same reasons as amended claim 29.

Claim 48

Applicant has amended claim 48 to recited “an input/output bus that is coupled to the computer-memory complex, at least one host device” As discussed above with regards to claim 18, Judd et al. does not disclose the recited host if the Examiner cites hosts 80 and 81 as the recited node. Thus, amended claim 48 is patentable over Judd et al. for at least the same reasons as amended claim 18.

Claims 49 to 61

Claims 49 to 61 depend from amended claim 48 and are patentable over the cited references for at least the same reasons as amended claim 48.

CONCLUSION

Applicant respectfully submits the Examiner has failed to show that the cited references disclose all the recited elements of claims 9, 13, 16, 18, 19, 21 to 25, 29, and 35 to 61. Accordingly, Applicant requests the rejections of claims 9, 13, 16, 18, 19, 21 to 25, 29, and 35 to 61 to be reversed.

I hereby certify that this correspondence is being transmitted prior to expiration of the set period of time by being transmitted via the Office electronic filing system in accordance with § 1.6(a)(4).

/David C Hsia/
Signature

September 17, 2008
Date

Respectfully submitted,

/David C Hsia/

David C. Hsia
Attorney for Applicant(s)
Reg. No. 46,235

Patent Law Group LLP
2635 North First St., Ste. 223
San Jose, California 95134
408-382-0480x206

CLAIMS APPENDIX

Claim 9: The node controller of claim 18, further comprising:

command queues coupled to the logic engine, the command queues operable to store logic control blocks which can be processed by the logic engine.

Claim 13: The node controller of Claim 18, wherein the node controller is implemented as an integrated circuit device.

Claim 16: The node controller of Claim 9, further comprising:

a producer register operable to specify a first address of a command queue; and

a consumer register operable to specify a second address of a command queue.

Claim 18: A node controller for a node in a data storage system having at least two nodes, each node comprising one computer-memory complex and one node controller distinct from said one computer-memory complex, the node controller comprising:

a memory controller operable to interface the node controller with (1) a cluster memory that stores data being transferred through the node, and (2) a link coupled to another node controller in another node of the data storage system;

an input/output bus interface operable to interface the node controller with an input/output bus coupled to a computer-memory complex of the node, at least one host device and at least one storage device;

a logic engine coupled to (1) the memory controller, (2) the link, and (3) the input/output bus interface;

wherein in a first type of data transfer, the logic engine performs a logic operation to a plurality of data from one of a plurality of data sources in the data storage system and writes the result of the logic operation to one of a plurality of data destinations in the data storage system, the data sources comprising the cluster memory and the input/output bus interface, the data destinations comprising the cluster memory, the link, and the input/output bus interface, the logic operation being used to calculate a parity data for writing a full or a partial RAID stripe or to reconstruct a lost data using the parity data.

Claim 19: The node controller of claim 18, wherein in a second type of data transfer, the input/output bus interface writes a data into the cluster memory and in response the logic engine copies the data to the another node via the link.

Claim 21: The node controller of claim 18, wherein the input/output bus interface comprises a peripheral component interconnect (PCI) control interface and the input/output bus comprises a PCI bus.

Claim 22: The node controller of claim 21, wherein the computer-memory complex manages the PCI bus.

Claim 23: The node controller of claim 22, wherein the computer-memory complex supports a service selected from the group consisting of a HTTP service, a NFS service, and a CIFS service.

Claim 24: The node controller of claim 18, wherein the computer-memory complex is not burdened with temporarily storing data being transferred through the node in the computer-memory complex.

Claim 25: The node controller of claim 18, wherein the logic operation comprises an XOR operation.

Claim 29: A node in a data storage system comprising at least two nodes, the node comprising:

- an input/output bus coupled to at least one host and at least one storage device;

- a computer-memory complex, comprising:

- a central processing unit (CPU);

- a system memory storing information for controlling data transfer through the node;
 - and

- a controller coupling the CPU, the system memory, and the input/output bus; and

- a node controller distinct from the computer-memory complex, the node controller comprising:

- a memory controller operable to interface the node controller with a cluster memory of the node, the cluster memory storing data being transferred through the node;

an input/output bus interface operable to interface the node controller with the input/output bus;

a logic engine coupled to the memory controller, the input/output bus interface, and a link to another node of the data storage system;

wherein at least one of the host device and the storage device coupled to the input/output bus is able to read and write the cluster memory via the input/output bus, the logic engine is able to transfer data from the cluster memory to the another node via the link, and the memory controller is able to receive data from the another node via the link.

Claim 35: The node of claim 29, wherein in a data transfer, the logic engine performs a logic operation on data from at least one of the input/output bus and the cluster memory, and writes a result of the logic operation to at least one of the input/output bus, the cluster memory, and the link.

Claim 36: The node of claim 35, further comprising a command queue operable to store a logic control block to be processed by the logic engine, the logic control block specifying at least one data source and at least one data destination for the logic operation.

Claim 37: The node of claim 36, further comprising:

a producer register operable to specify a first address of the command queue; and

a consumer register operable to specify a second address of the command queue.

Claim 38: The node of claim 35, wherein the logic engine comprises an exclusive OR (XOR) engine.

Claim 39: The node of claim 38, wherein the XOR engine is used to calculate a parity data for writing a full or a partial RAID stripe.

Claim 40: The node of claim 38, wherein the XOR engine is used to reconstruct a lost data using a parity data.

Claim 41: The node of claim 29, wherein the node controller is implemented as an integrated circuit device.

Claim 42: The node of claim 29, wherein the input/output bus interface comprises a peripheral component interconnect (PCI) control interface and the input/output bus comprises a PCI bus.

Claim 43: The node of claim 29, wherein the node controller is operable to be programmed by the computer-memory complex.

Claim 44: The node of claim 29, wherein in a data transfer, the input/output bus interface writes a data into the cluster memory and in response the logic engine copies the data to the another node via the link.

Claim 45: The node of claim 29, wherein the computer-memory complex supports a service selected from the group consisting of a HTTP service, a NFS service, and a CIFS service.

Claim 46: The node of claim 29, wherein the computer-memory complex is not burdened with temporarily storing data being transferred through the node in the system memory.

Claim 47: The node of claim 29, further comprising:

another input/output bus interface operable to interface the node controller with another input/output bus coupled to the computer-memory complex and at least one of another host device and another storage device, wherein the computer-memory complex, the another host device, and the another storage device are able to read and write the cluster memory via the another input/output bus.

Claim 48: A node in a data storage system having at least one node for providing access to a data storage facility, the node comprising a node controller and a computer memory complex, the node controller distinct from the computer-memory complex, the node controller comprising:

an input/output bus interface operable to interface the node controller with an input/output bus that is coupled to the computer-memory complex, at least one host device, and at least one storage device;

a memory controller operable to interface the node controller with a cache memory in the node;

one or more logic engines coupled to the input/output bus interface, to the memory controller, and to a link that can be coupled to another node in the data storage system;

wherein the node controller is arranged so that the computer-memory complex, the host device, and the storage device are able to access the node controller via the input/output bus so that the host device and the storage device are able to read and write the cache memory via the input/output bus; and

wherein the logic engine is able to transfer data from the cache memory to another node in the data storage system via the link, and the memory controller is able to receive data from the another node in the data storage system via the link.

Claim 49: The node of claim 48, wherein the logic engine is operable to perform a logic operation in a data transfer on data from at least one of the input/output bus and the cluster memory, and to write a result of the logic operation to at least one of the input/output bus, the cluster memory and the link.

Claim 50: The node of claim 49, wherein the logic engine comprises an exclusive OR engine.

Claim 51: The node of claim 49, the node controller further comprising a command queue operable to store a logic control block to be processed by the logic engine, the logic control block specifying a data source and data destination for the logic operation.

Claim 52: The node of claim 48, wherein the node controller is implemented as an integrated circuit device.

Claim 53: The node of claim 48, wherein the input/output bus interface comprises a peripheral component interconnect control interface and the input/output bus comprises a peripheral component interconnect bus.

Claim 54: The node of claim 48, wherein the node controller is operable to be programmed by the computer-memory complex.

Claim 55: The node of claim 51, the node controller further comprising:

a producer register operable to specify a first address of the command queue; and

a consumer register operable to specify a second address of the command queue.

Claim 56: The node of claim 48, wherein the node controller is arranged to send data written into the cache memory to the another node via the link.

Claim 57: The node of claim 48, wherein the computer-memory complex is arranged to support a service selected from the group consisting of a HTTP service, a NFS service, and a CIFS service.

Claim 58: The node of claim 48, wherein the computer-memory complex is not burdened with temporarily storing data being transferred through the node in the computer-memory complex.

Claim 59: The node of claim 50, wherein the exclusive OR engine is arranged to calculate a parity data for writing a full or a partial RAID stripe.

Claim 60: The node of claim 50, wherein the exclusive OR engine is arranged to reconstruct lost data using a parity data.

Claim 61: The node of claim 48, wherein the node controller is configured to act as a slave device.

EVIDENCE APPENDIX

None

RELATED PROCEEDINGS APPENDIX

None.